

# SURVEILLANCE ON THE FLY

**C**AMERAS mounted on an airplane or other airborne platform constantly monitor an area, feeding data to the ground for real-time analysis. This R&D 100 Award-winning technology, called the Sonoma Persistent Surveillance System, can provide continuous, real-time video imagery of an area the size of a small city with a resolution fine enough to track 8,000 moving objects in its field of view.

With Sonoma, a user can ascertain the pattern of movements, for example, of vehicles, both spatially and over time. This capability is essential for establishing connections between known and unknown targets, determining the time history of those connections, and identifying new targets. Sonoma was developed for nonproliferation applications by researchers in Livermore's Nonproliferation, Homeland and International Security Directorate.

But it can also provide real-time data to monitor traffic, enhance security at special events, and improve surveillance at borders and ports. Livermore physicist Deanna Pennington, who leads the development project, says, "If a Sonoma system had been used in the aftermath of the Katrina and Rita hurricanes, emergency responders would have had real-time information on roads, water levels, and traffic conditions, which perhaps would have saved lives."

## End-to-End Systems Approach

The key technology developments that made Sonoma possible include novel sensor designs that can view a wide area at high resolution, real-time on-board data processing, automated data processing to absorb and analyze the resulting data stream, and high-performance visualization technology. Livermore's end-to-end systems



Boxes superimposed on a satellite image of Washington, DC, show the area coverage provided by a standard video camera, a high-definition television (HDTV) camera, a 22-megapixel sensor (Sonoma prototype 1), a 66-megapixel camera (prototype 2), and a 176-megapixel camera (prototype 3). The image in the lower right corner indicates the resolution provided by the Sonoma sensors.



The Sonoma development team (from left to right): Michael Newman, David Bloom, Charles Thompson, Deanna Pennington, Curtis Brown, Aaron Wegner, Laurence Flath, Robert Sawvel, and Michael Kartz. Not pictured: Allen House, Daniel Knight, John Marion, and Gary Stone.

Lawrence Livermore National Laboratory



approach to problem-solving and expertise in optics and advanced computing fostered the needed developments in sensor and data-processing technologies.

For the sensor, the Sonoma team designed a mosaic technique that optically stitches together images from multiple cameras to create a large-format sensor. The technique can be scaled to produce large-aperture arrays that image in both the visible and infrared spectra. It requires far less computing power than other stitching methods, so computer resources are available for other data-processing routines. Another advantage is that the sensor's commercially available cameras, lenses, and mounting hardware cost just one-tenth of those used in similar surveillance sensors.

The top figure on p. 4 compares the area coverage provided by various sensors superimposed on a satellite image of Washington, DC. The current Sonoma prototype, a 66-megapixel sensor, can cover the entire central urban area with an image resolution comparable to that shown in the lower right corner.

A 176-megapixel sensor (prototype 3) is under construction.

For surveillance data to be useful, an analyst must be able to see the information in real time. In other words, all data processing for one frame must be completed before the next frame is captured. With data being collected at two frames per second, the data generated exceed the bandwidth of commercially available communications links by a factor of 100 to 10,000. Compressing the data before transmitting them to a ground-based station would cause artifacts, or errors, that can lead to misinterpretations. To eliminate this problem, the Sonoma team developed a data-processing system for the airborne platform based on the graphics processors used in gaming applications.

The software and hardware for the onboard system work extremely fast. They collect and archive the video data, record an object's geospatial coordinates to freeze the background, and change the viewing perspective so that data collected at a 30- to 45-degree angle are displayed from the top down. All of these data are processed at a rate of two frames per second.

### Freezing the Background

Raw imagery collected from a moving platform can be difficult to interpret because its perspective changes constantly as the platform moves. Removing this motion and registering an image's geospatial coordinates in the traditional manner require either data from multiple Global Positioning System (GPS) satellites or scene-based correlation algorithms, both of which are computationally intensive and time consuming.

The Livermore team solved this problem by combining a GPS/inertial measurement unit (IMU) with the sensor to improve tracking accuracy. The GPS/IMU and sensor operate through

a common boresight, and both are mounted on a gimbal, which allows them to rotate in three dimensions. The center of the image remains locked on target as the platform moves. Geospatial coordinates and elevation data are recorded into the header of each camera frame, so every pixel in every frame can be located accurately. Location data are then fed into a flat-earth model. Together, these operations subtract the motion of the platform so that moving objects can be easily separated from the background. Specialized software operates on graphics processing units rather than on standard computer processing units to speed the processing of visual data. (See *S&TR*, November 2005, pp. 19–20.)

### The System at Work

Data processed on the airborne platform are handled in two ways. Full-resolution data on the regions around each moving object are transmitted to a ground-based station twice per second via a radio-frequency link. Full-resolution contextual background imagery is sent to the station once per minute, where it is combined with the moving object data for real-time display. The resulting reduction in data volume is equivalent to a compression of up to 100 million but without the artifacts.

The Sonoma team developed several visualization tools to give the ground-based analysts a range of options to manipulate the incoming data. Users can view data from all of the cameras in the sensor array or just one, panning over an area or zooming in and out. They also can alter the rate that imagery is played forward or backward, change the scale for an area of interest, or dynamically reset the magnification. Another visualization tool builds the merged, aggregate image file, using data from all of the cameras or only a subset.

Together, these components—the broad-area imager, onboard data acquisition and processing, a communications link, and a ground station to reconstruct and analyze the imagery in real time—make Sonoma a complete end-to-end system. Other sensor packages or custom image-processing techniques are available, but Sonoma is the only fully integrated system available. Says Pennington, “Our success with Sonoma is already changing how surveillance systems are designed, in terms of a system's architecture and the concepts of operation.”

—Katie Walter

**Key Words:** electro-optical sensors, graphics processing unit, nonproliferation, R&D 100 Award, Sonoma Persistent Surveillance System, video imagery, visualization.

**For further information contact Deanna Pennington (925) 423-9234 (pennington1@llnl.gov).**